Reconfiguration on sparse graphs

Daniel Lokshtanov[1], Amer E. Mouawad[2], Fahad Panolan[3], M.S. Ramanujan[1], and Saket Saurabh[3]

[1] University of Bergen, Norway.
`daniello,ramanujan.sridharan@ii.uib.no`
[2] David R. Cheriton School of Computer Science
University of Waterloo, Ontario, Canada.
`aabdomou@uwaterloo.ca`
[3] Institute of Mathematical Sciences
Chennai, India.
`fahad,saket@imsc.res.in`

Abstract. A vertex-subset graph problem $\mathcal{Q}$ defines which subsets of the vertices of an input graph are feasible solutions. A reconfiguration variant of a vertex-subset problem asks, given two feasible solutions S_s and S_t of size k, whether it is possible to transform S_s into S_t by a sequence of vertex additions and deletions such that each intermediate set is also a feasible solution of size bounded by k. We study reconfiguration variants of two classical vertex-subset problems, namely INDEPENDENT SET and DOMINATING SET. We denote the former by ISR and the latter by DSR. Both ISR and DSR are `PSPACE`-complete on graphs of bounded bandwidth and `W[1]`-hard parameterized by k on general graphs. We show that ISR is fixed-parameter tractable parameterized by k when the input graph is of bounded degeneracy or nowhere-dense. As a corollary, we answer positively an open question concerning the parameterized complexity of the problem on graphs of bounded treewidth. Moreover, our techniques generalize recent results showing that ISR is fixed-parameter tractable on planar graphs and graphs of bounded degree. For DSR, we show the problem fixed-parameter tractable parameterized by k when the input graph does not contain large bicliques, a class of graphs which includes graphs of bounded degeneracy and nowhere-dense graphs.

1 Introduction

Given an n-vertex graph G and two vertices s and t in G, determining whether there exists a path and computing the length of the shortest path between s and t are two of the most fundamental graph problems. In the classical battle of `P` versus `NP` or "easy" versus "hard", both of these problems are on the easy side. That is, they can be solved in $poly(n)$ time, where $poly$ is any polynomial function. But what if our input consisted of a 2^n-vertex graph? Of course, we can no longer assume G to be part of the input, as reading the input alone requires more than $poly(n)$ time. Instead, we are given an oracle encoded using $poly(n)$ bits and that can, in constant or $poly(n)$ time, answer queries of the form "is u a vertex in G" or "is there an edge between u and v?". Given such an oracle and two vertices of the 2^n-vertex graph, can we still determine if there is a path or compute the length of the shortest path between s and t in $poly(n)$ time?

A slightly different, but equally insightful, formulation of the question above is as follows. Given a set S of n objects, consider the graph $R(S)$ which contains one node for each set in the power set of S, 2^S, and two nodes are adjacent in $R(S)$ whenever the size of their symmetric difference is equal to one. Clearly, this graph contains 2^n nodes and can be easily encoded in $poly(n)$ bits using the oracle described above. It is not hard to see that there exists a path between any two nodes of $R(S)$. Moreover, computing the length of a shortest path can be accomplished in constant time; it is equal to the size of the symmetric difference of the two underlying sets. If the node set of $R(S)$ were instead restricted to a subset of 2^S, both of our problems can become `NP`-complete or even `PSPACE`-complete. Therefore, another interesting question is whether we can determine what types of "restriction" on the node set of $R(S)$ induce such variations in the complexity of the two problems.

These two seemingly artificial questions are in fact quite natural and appear in many practical and theoretical problems. In particular, these are exactly the types of questions asked under the reconfiguration

framework, the main subject of this work. Under the reconfiguration framework, instead of finding a feasible solution to some instance $\mathcal{I}$ of a search problem $\mathcal{Q}$, we are interested in structural and algorithmic questions related to the solution space of $\mathcal{Q}$. Naturally, given some adjacency relation $\mathcal{A}$ defined over feasible solutions of $\mathcal{Q}$, size of the symmetric difference being one such relation, the solution space can be represented using a graph $R_\mathcal{Q}(\mathcal{I})$. $R_\mathcal{Q}(\mathcal{I})$ contains one node for each feasible solution of $\mathcal{Q}$ on instance $\mathcal{I}$ and two nodes share an edge whenever their corresponding solutions are adjacent under $\mathcal{A}$. An edge in $R_\mathcal{Q}(\mathcal{I})$ corresponds to a *reconfiguration step*, a walk in $R_\mathcal{Q}(\mathcal{I})$ is a sequence of such steps, a *reconfiguration sequence*, and $R_\mathcal{Q}(\mathcal{I})$ is a *reconfiguration graph*.

Studying problems related to reconfiguration graphs has received considerable attention in recent literature [4, 22, 25, 26, 30, 34], the most popular problem being to determine whether there exists a reconfiguration sequence between two given feasible solution. In most cases, this problem was shown PSPACE-hard in general, although some polynomial-time solvable restricted cases have been identified. For PSPACE-hard cases, it is not surprising that shortest paths between solutions can have exponential length. More surprising is that for most known polynomial-time solvable cases the diameter of the reconfiguration graph has been shown to be polynomial. Some of the problems that have been studied under the reconfiguration framework include INDEPENDENT SET [31], VERTEX COVER [33], SHORTEST PATH [5, 30], COLORING [3, 6, 7, 9–11, 29], and BOOLEAN SATISFIABILITY [22]. We refer the reader to the recent survey by Van den Heuvel [43] for a detailed overview. Recently, a systematic study of the parameterized complexity of reconfiguration problems was initiated by Mouawad et al. [34]; various problems were identified where the problem was not only NP-hard (or PSPACE-hard), but also W-hard under various parameterizations.

Overview of our results. In this work, we focus on reconfiguration variants of the INDEPENDENT SET (IS) and DOMINATING SET (DS) problems. Given two independent sets I_s and I_t of a graph G such that $|I_s| = |I_t| = k$, the INDEPENDENT SET RECONFIGURATION problem asks whether there exists a sequence of independents sets $\sigma = \langle I_0, I_1, \ldots, I_\ell \rangle$, for some ℓ, such that:

(1) $I_0 = I_s$ and $I_\ell = I_t$,
(2) I_i is an independent set of G for all $0 \le i \le \ell$,
(3) $|I_i \Delta I_{i+1}| = 1$ for all $0 \le i < \ell$, and
(4) $k - 1 \le |S_i| \le k$ for all $0 \le i \le \ell$.

Alternatively, given a graph G and integer k, the reconfiguration graph $R_{\mathrm{IS}}(G, k-1, k)$ has a node for each independent set of G of size k or $k - 1$ and two nodes are adjacent in $R_{\mathrm{IS}}(G, k-1, k)$ whenever the corresponding independent sets can be obtained from one another by either the addition or the deletion of a single vertex. The reconfiguration graph $R_{\mathrm{DS}}(G, k, k+1)$ is defined similarly for dominating sets. Hence, ISR and DSR can be formally stated as follows:

INDEPENDENT SET RECONFIGURATION (ISR)
Input: Graph G, positive integer k, and two k-independent sets I_s and I_t
Question: Is there a path from I_s to I_t in $R_{\mathrm{IS}}(G, k-1, k)$?

DOMINATING SET RECONFIGURATION (DSR)
Input: Graph G, positive integer k, and two k-dominating sets D_s and D_t
Question: Is there a path from D_s to D_t in $R_{\mathrm{DS}}(G, k, k+1)$?

Note that since we only allow independent sets of size k and $k - 1$ the ISR problem is equivalent to reconfiguration under the token jumping model considered by Ito et al. [27, 28]. ISR is known to be PSPACE-complete on graphs of bounded bandwidth [35, 44] (hence pathwidth and treewidth) and W[1]-hard on general graphs [28]. On the positive side, the problem was shown fixed-parameter tractable, with parameter k, for graphs of bounded degree, planar graphs, and graphs excluding $K_{3,d}$ as a (not necessarily induced) subgraph, for any constant d [27, 28]. We push this boundary further by showing that the problem remains fixed-parameter tractable for graphs of bounded degeneracy and nowhere-dense graphs (Figure 1). As a

corollary, we answer positively an open question concerning the parameterized complexity of the problem (parameterized by k) on graphs of bounded treewidth.

For DSR, we first show that the problem is `W[1]`-hard on general graphs by adapting the well-known (parameter-preserving) reduction from INDEPENDENT SET to DOMINATING SET. Then, we show that the problem is fixed-parameter tractable, with parameter k, for graphs excluding $K_{d,d}$ as a (not necessarily induced) subgraph, for any constant d. Note that this class of graphs includes both nowhere-dense and bounded degeneracy graphs and is the "largest" class on which the DOMINATING SET problem is known to be in `FPT` [40, 42].

Clearly, our main open question is whether ISR remains fixed-parameter tractable on graphs excluding $K_{d,d}$ as a subgraph. Intuitively, all of the classes we consider fall under the category of "sparse" graph classes. Hence, in some sense, one would not expect a sparse graph to have "too many" dominating sets of fixed small size k as n becomes larger and larger. For independent sets, the situation is reversed. As n grows larger, so does the number of independent sets of fixed size k. So it remains to be seen whether some structural properties of graphs excluding $K_{d,d}$ as a subgraph can be used to settle our open question or whether the problem becomes `W[1]`-hard. In the latter case, this would be the first example of a `W[1]`-hard problem (in general), which is in `FPT` on a class $\mathcal{C}$ of graphs but where the reconfiguration version is not; finding such a problem, we believe, is interesting in its own right. Another open question is whether we can adapt our results for ISR to find shortest reconfiguration sequences. Our algorithm for DSR does in fact guarantee shortest reconfiguration sequences but, as we shall see, the same does not hold for both ISR algorithms.

2   Preliminaries

For an in-depth review of general graph theoretic definitions we refer the reader to the book of Diestel [16]. Unless otherwise stated, we assume that each graph G is a simple, undirected graph with vertex set $V(G)$ and edge set $E(G)$, where $|V(G)| = n$ and $|E(G)| = m$. The *open neighborhood*, or simply *neighborhood*, of a vertex v is denoted by $N_G(v) = \{u \mid uv \in E(G)\}$, the *closed neighborhood* by $N_G[v] = N_G(v) \cup \{v\}$. Similarly, for a set of vertices $S \subseteq V(G)$, we define $N_G(S) = \{v \mid uv \in E(G), u \in S, v \notin S\}$ and $N_G[S] = N_G(S) \cup S$. The *degree* of a vertex is $|N_G(v)|$. We drop the subscript G when clear from context. A *subgraph* of G is a graph G' such that $V(G') \subseteq V(G)$ and $E(G') \subseteq E(G)$. The *induced subgraph* of G with respect to $S \subseteq V(G)$ is denoted by $G[S]$; $G[S]$ has vertex set S and edge set $\{uv \in E(G[S]) \mid u, v \in S, uv \in E(G)\}$. We denote by $\Delta(G)$ and $\delta(G)$ the maximum and minimum degree of G, respectively.

A *walk* of length ℓ from v_0 to v_ℓ in G is a vertex sequence $v_0, \ldots, v_\ell$, such that for all $i \in \{0, \ldots, \ell - 1\}$, $v_i v_{i+1} \in E(G)$. It is a *path* if all vertices are distinct. It is a *cycle* if $\ell \geq 3$, $v_0 = v_\ell$, and $v_0, \ldots, v_{\ell-1}$ is a path. A path from vertex u to vertex v is also called a *uv-path*. The *distance* between two vertices u and v of G, $dist_G(u, v)$, is the length of a shortest uv-path in G (positive infinity if no such path exists). The *eccentricity* of a vertex $v \in V(G)$, $ecc(v)$, is equal to $max_{u \in V(G)}(dist_G(u, v))$. The *radius* of G, $rad(G)$, is equal to $min_{v \in V(G)}(ecc(v))$. The *diameter* of G, $diam(G)$, is equal to $max_{v \in V(G)}(ecc(v))$. For $r \geq 0$, the *r-neighborhood* of a vertex $v \in V(G)$ is defined as $N_G^r[v] = \{u \mid dist_G(u, v) \leq r\}$. We write $B(v, r) = N_G^r[v]$ and call it a *ball of radius r around v*; for $S \subseteq V(G)$, $B(S, r) = \bigcup_{v \in S} N_G^r[v]$.

Contracting an edge uv of G results in a new graph H in which the vertices u and v are deleted and replaced by a new vertex w that is adjacent to $N_G(u) \cup N_G(v) \setminus \{u, v\}$. If a graph H can be obtained from G by repeatedly contracting edges, H is said to be a *contraction* of G. If H is a subgraph of a contraction of G, then H is said to be a *minor* of G, denoted by $H \preceq_m G$. An equivalent characterization of minors states that H is a minor of G if there is a map that associates to each vertex v of H a non-empty connected subgraph G_v of G such that G_u and G_v are disjoint for $u \neq v$ and whenever there is an edge between u and v in H there is an edge in G between some node in G_u and some node in G_v. The subgraphs G_v are called *branch sets*. H is a *minor at depth r of G*, $H \preceq_m^r G$, if H is a minor of G which is witnessed by a collection of branch sets $\{G_v \mid v \in V(H)\}$, each of which induces a graph of radius at most r. That is, for each $v \in V(H)$, there is a $w \in V(G_v)$ such that $V(G_v) \subseteq N_{G_v}^r[w]$.

3

Sparse graph classes. We define the three main classes we consider. Figure 1 illustrates the relationship between these classes and some other well-known classes of sparse graphs. We refer the reader to [8, 38, 36] for more details.

Definition 1 ([38, 36]). *A class of graphs* $\mathcal{C}$ *is said to be* nowhere-dense *if for every* $d \geq 0$ *there exists a graph* H_d *such that* $H_d \npreceq_m^d G$ *for all* $G \in \mathcal{C}$. $\mathcal{C}$ *is* effectively nowhere-dense *if the map* $d \mapsto H_d$ *is computable. Otherwise,* $\mathcal{C}$ *is said to be* somewhere-dense.

Nowhere-dense classes of graphs were introduced by Nesetril and Ossona de Mendez [38, 36] and "nowhere-density" turns out to be a very robust concept with several natural characterizations [23]. We use one such characterization in Section 3.2. It follows from the definition that planar graphs, graphs of bounded treewidth, graphs of bounded degree, H-minor-free graphs, and H-topological-minor-free graphs are nowhere-dense [38, 36]. As in the work of Dawar and Kreutzer [14], we are only interested in effectively nowhere-dense classes; all natural nowhere-dense classes are effectively nowhere-dense, but it is possible to construct artificial classes that are nowhere-dense, but not effectively so.

Definition 2. *A class of graphs* $\mathcal{C}$ *is said to be* d-degenerate *if there is an integer* d *such that every induced subgraph of any graph* $G \in \mathcal{C}$ *has a vertex of degree at most* d.

Graphs of bounded degeneracy and nowhere-dense graphs are incomparable [24]. In other words, graphs of bounded degeneracy are somewhere-dense.

Proposition 1 ([32]). *The number of edges in a* d-degenerate graph is at most dn and hence its average degree is at most $2d$.

Degeneracy is a hereditary property, hence any induced subgraph of a d-degenerate graph is also d-degenerate. It is well-known that graphs of treewidth at most d are also d-degenerate. Moreover a d-degenerate graph cannot contain $K_{d+1,d+1}$ as a subgraph, which brings us to the class of biclique-free graphs. The relationship between bounded degeneracy, nowhere-dense, and $K_{d,d}$-free graphs was shown by Philip et al. and Telle and Villanger [40, 42].

Definition 3. *A class of graphs* $\mathcal{C}$ *is said to be* d-biclique-free, *for some* $d > 0$, *if* $K_{d,d}$ *is not a subgraph of any* $G \in \mathcal{C}$, *and it is said to be* biclique-free *if it is* d-biclique-free *for some* d.

Proposition 2 ([40, 42]). *Any degenerate or nowhere-dense class of graphs is biclique-free, but not vice-versa.*
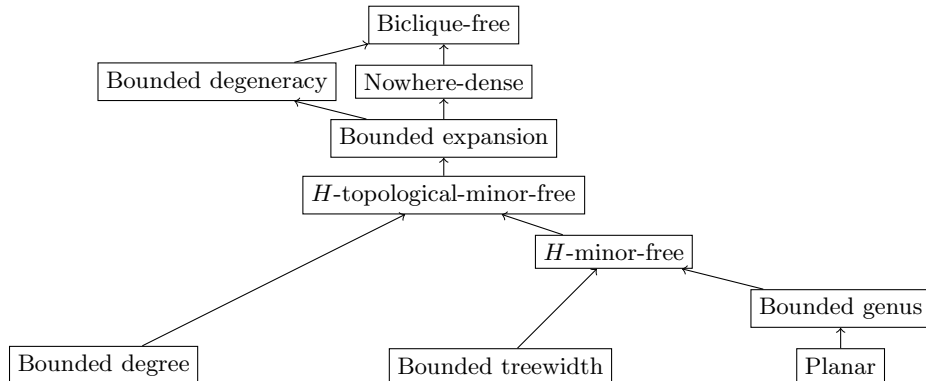

Fig. 1. Sparse graph classes [8, 38, 36]. Arrows indicate inclusion.

Parameterized complexity. Using the framework developed by Downey and Fellows [17], a *parameterized problem* includes in the input a parameter p. For a parameterized problem $\mathcal{Q}$ with inputs of the form (x, p), $|x| = n$ and p a positive integer, $\mathcal{Q}$ is *fixed-parameter tractable* (or in FPT) if it can be decided in $f(p)n^c$ time, where f is an arbitrary function and c is a constant independent of both n and p. $\mathcal{Q}$ is in the class XP if it can be decided in $n^{f(p)}$ time. $\mathcal{Q}$ has a *kernel* of size $f(p)$ if there is an algorithm that transforms the input (x, p) to (x', p') in polynomial time (with respect to $|x|$ and p) such that (x, p) is a yes-instance if and only if (x', p') is a yes-instance, $p' \leq g(p)$, and $|x'| \leq f(p)$. Each problem in FPT has a kernel, possibly of exponential (or worse) size [17].

In order to distinguish between parameterized problems solvable in $n^{f(p)}$ time and parameterized problems solvable in $f(p)n^c$ time, Downey and Fellows [17] introduced the W-*hierarchy*. The hierarchy consists of a complexity class W[t] for every integer $t \geq 1$ such that W[t] $\subseteq$ W[t + 1] for all t. They proved that FPT $\subseteq$ W[1] $\subseteq$ W[2] $\subseteq \ldots \subseteq$ W[t] and conjectured that strict containment holds. In particular, the assumption FPT $\subset$ W[1] is a natural parameterized analogue of the conjecture that P $\neq$ NP. Moreover, Downey and Fellows showed that the INDEPENDENT SET problem parameterized by solution size is W[1]-complete and the DOMINATING SET problem parameterized by solution size is W[2]-complete. Showing hardness in the parameterized setting is usually accomplished using FPT reductions. The reader is referred to the books of Niedermeier, Flum, and Grohe for more on parameterized complexity [21, 39].

Reconfiguration. For any vertex-subset problem $\mathcal{Q}$, graph G, and positive integer k, we consider the *reconfiguration graph* $R_{\mathcal{Q}}(G, k, k + 1)$ when $\mathcal{Q}$ is a minimization problem (e.g. DOMINATING SET) and the reconfiguration graph $R_{\mathcal{Q}}(G, k - 1, k)$ when $\mathcal{Q}$ is a maximization problem (e.g. INDEPENDENT SET). A set $S \subseteq V(G)$ has a corresponding node in $V(R_{\mathcal{Q}}(G, r_l, r_u))$, $r_l \in \{k-1, k\}$ and $r_u \in \{k, k+1\}$, if and only if S is a feasible solution for $\mathcal{Q}$ and $r_l \leq |S| \leq r_u$. We refer to *vertices* in G using lower case letters (e.g. u, v) and to the *nodes* in $R_{\mathcal{Q}}(G, r_l, r_u)$, and by extension their associated feasible solutions, using upper case letters (e.g. A, B). If $A, B \in V(R_{\mathcal{Q}}(G, r_l, r_u))$ then there exists an edge between A and B in $R_{\mathcal{Q}}(G, r_l, r_u)$ if and only if there exists a vertex $u \in V(G)$ such that $\{A \setminus B\} \cup \{B \setminus A\} = \{u\}$. Equivalently, for $A \Delta B = \{A \setminus B\} \cup \{B \setminus A\}$ the *symmetric difference* of A and B, A and B share an edge in $R_{\mathcal{Q}}(G, r_l, r_u)$ if and only if $|A \Delta B| = 1$.

We write $A \leftrightarrow B$ if there exists a path in $R_{\mathcal{Q}}(G, r_l, r_u)$, a reconfiguration sequence, joining A and B. Any reconfiguration sequence from *source* feasible solution S_s to *target* feasible solution S_t, which we sometimes denote by $\sigma = \langle S_0, S_1, \ldots, S_\ell \rangle$, for some ℓ, has the following properties:

- $S_0 = S_s$ and $S_\ell = S_t$,
- S_i is a feasible solution for $\mathcal{Q}$ for all $0 \leq i \leq \ell$,
- $|S_i \Delta S_{i+1}| = 1$ for all $0 \leq i < \ell$, and
- $r_l \leq |S_i| \leq r_u$ for all $0 \leq i \leq \ell$.

We denote the *length* of σ by $|\sigma|$. For $0 < i \leq \ell$, we say vertex $v \in V(G)$ is *added* at step/index/position/slot i if $v \notin S_{i-1}$ and $v \in S_i$. Similarly, a vertex v is *removed* at step/index/position/slot i if $v \in S_{i-1}$ and $v \notin S_i$. A vertex $v \in V(G)$ is *touched* in the course of a reconfiguration sequence if v is either added or removed at least once; it is *untouched* otherwise. A vertex is *removable* (*addable*) from feasible solution S if $S \setminus \{v\}$ ($S \cup \{v\}$) is also a feasible solution for $\mathcal{Q}$. For any pair of consecutive solutions (S_{i-1}, S_i) in σ, we say S_i (S_{i-1}) is the *successor* (*predecessor*) of S_{i-1} (S_i). A reconfiguration sequence $\sigma' = \langle S_0, S_1, \ldots, S_{\ell'} \rangle$ is a *prefix* of $\sigma = \langle S_0, S_1, \ldots, S_\ell \rangle$ if $\ell' < \ell$.

We adapt the concept of irrelevant vertices from parameterized complexity to introduce the notions of irrelevant and strongly irrelevant vertices for reconfiguration. Since these notions apply to almost any reconfiguration problem, we give general definitions.

Definition 4. *For any vertex-subset problem $\mathcal{Q}$, n-vertex graph G, positive integers r_l and r_u, and $S_s, S_t \in V(R_{\mathcal{Q}}(G, r_l, r_u))$ such that there exists a reconfiguration sequence from S_s to S_t in $R_{\mathcal{Q}}(G, r_l, r_u)$, we say a vertex $v \in V(G)$ is* irrelevant *(with respect to S_s and S_t) if and only if $v \notin S_s \cup S_t$ and there exists a reconfiguration sequence from S_s to S_t in $R_{\mathcal{Q}}(G, r_l, r_u)$ which does not touch v. We say v is* strongly irrelevant *(with respect to S_s and S_t) if it is irrelevant and the length of a shortest reconfiguration sequence*

from S_s to S_t which does not touch v is no greater than the length of a shortest reconfiguration sequence which does (if the latter sequence exists).

At a high level, it is enough to consider irrelevant vertices when trying to find *any* reconfiguration sequence between two feasible solutions, but strongly irrelevant vertices must be considered if we wish to find a *shortest* reconfiguration sequence. As we shall see, our algorithm for DSR does in fact find strongly irrelevant vertices and can therefore be used to find shortest reconfiguration sequences. For ISR, we are only able to find irrelevant vertices and reconfiguration sequences are not guaranteed to be of shortest possible length.

3 Independent set reconfiguration

3.1 Graphs of bounded degeneracy

To show that the ISR problem is fixed-parameter tractable on d-degenerate graphs, for some integer d, we will proceed in two stages. In the first stage, we will show, for an instance (G, I_s, I_t, k), that as long as the number of low-degree vertices in G is "large enough" we can find an irrelevant vertex (Definition 4). Once the number of low-degree vertices is bounded, a simple counting argument (Proposition 3) shows that the size of the remaining graph is also bounded and hence we can solve the instance by exhaustive enumeration.

Proposition 3. *Let G be an n-vertex d-degenerate graph, $S_1 \subseteq V(G)$ be the set of vertices of degree at most $2d$, and $S_2 = V(G) \setminus S_1$. If $|S_1| < s$, then $|V(G)| \leq (2d+1)s$.*

Proof. The number of edges in a d-degenerate graph is at most dn and hence its average degree is at most $2d$ (Proposition 1). If $|V(G)| = (2d+1)s + c$, for $c \geq 1$, then $|S_2| = |V(G) \setminus S_1| > 2ds + c$, $\sum_{v \in S_2} |N_G(v)| > (2ds+c)(2d+1)$, and we obtain the following contradiction:

$$\frac{\sum_{v \in S_1} |N_G(v)| + \sum_{v \in S_2} |N_G(v)|}{|V(G)|} > \frac{(2ds+c)(2d+1)}{(2d+1)s+c}$$
$$= \frac{4d^2 s + 2ds + 2dc + c}{(2d+1)s+c}$$
$$= \frac{2d(2ds+s+c)+c}{2ds+s+c} > 2d.$$

$\square$

To find irrelevant vertices, we make use of the following classical result of Erdõs and Rado [20], also known in the literature as the sunflower lemma. We first define the terminology used in the statement of the theorem. A *sunflower* with k *petals* and a *core* Y is a collection of sets $S_1, \ldots, S_k$ such that $S_i \cap S_j = Y$ for all $i \neq j$; the sets $S_i \setminus Y$ are petals and we require none of them to be empty. Note that a family of pairwise disjoint sets is a sunflower (with an empty core).

Theorem 1 (Sunflower Lemma [20]). *Let $\mathcal{A}$ be a family of sets (without duplicates) over a universe $\mathcal{U}$, such that each set in $\mathcal{A}$ has cardinality at most d. If $|\mathcal{A}| > d!(k-1)^d$, then $\mathcal{A}$ contains a sunflower with k petals and such a sunflower can be computed in time polynomial in $|\mathcal{A}|$, $|\mathcal{U}|$, and k.*

Lemma 1. *Let (G, I_s, I_t, k) be an instance of ISR where G is d-degenerate and let B be the set of vertices in $V(G) \setminus \{I_s \cup I_t\}$ of degree at most $2d$. If $|B| > (2d+1)!(2k-1)^{2d+1}$, then there exists an irrelevant vertex $v \in V(G) \setminus \{I_s \cup I_t\}$ such that (G, I_s, I_t, k) is a yes-instance if and only if (G', I_s, I_t, k) is a yes-instance, where G' is obtained from G by deleting v and all edges incident on v.*

Proof. Let $b_1, b_2, \ldots, b_{|B|}$ denote the vertices in B and let $\mathcal{A} = \{N_G[b_1], N_G[b_2], \ldots, N_G[b_{|B|}]\}$ denote the family of sets corresponding to the closed neighborhoods of each vertex in B and set $\mathcal{U} = \bigcup_{b \in B} N[b]$. Since

$|B|$ is greater than $(2d+1)!(2k-1)^{2d+1}$, we know from Theorem 1 that $\mathcal{A}$ contains a sunflower with $2k$ petals and such a sunflower can be computed in time polynomial in $|\mathcal{A}|$ and k. Note that we assume, without loss of generality, that there are no two vertices u and v in $V(G) \setminus \{I_s \cup I_t\}$ such that $N_G[u] = N_G[v]$, as we can safely delete one of them from the input graph otherwise, i.e. one of the two is (strongly) irrelevant. Let v_{ir} be a vertex whose closed neighborhood corresponds to one of those $2k$ petals. We claim that v_{ir} is irrelevant and can therefore be deleted from G to obtain G'.

To see why, consider any reconfiguration sequence $\sigma = \langle I_s = I_0, I_1, \ldots, I_t = I_\ell \rangle$ from I_s to I_t in $R_{\mathrm{IS}}(G, k-1, k)$. Since $v_{ir} \notin I_s \cup I_t$, we let p, $0 < p < \ell$, be the first index in σ at which v_{ir} is added, i.e. $v_{ir} \in I_p$ and $v_{ir} \notin I_i$ for all $i < p$. Moreover, we let $q + 1$, $p < q + 1 \leq \ell$ be the first index after p at which v_{ir} is removed, i.e. $v_{ir} \in I_q$ and $v_{ir} \notin I_{q+1}$. We will consider the subsequence $\sigma_s = \langle I_p, \ldots, I_q \rangle$ and show how to modify it so that it does not touch v_{ir}. Applying the same procedure to every such subsequence in σ suffices to prove the lemma.

Since the sunflower constructed to obtain v_{ir} has $2k$ petals and the size of any independent set in σ (or any reconfiguration sequence in general) is at most k, there must exist another *free* vertex v_{fr} whose closed neighborhood corresponds to one of the remaining $2k - 1$ petals which we can add at index p instead of v_{ir}, i.e. $v_{fr} \notin N_G[I_p]$. We say v_{fr} *represents* v_{ir}. Assume that no such vertex exists. Then we know that either some vertex in the core of the sunflower is in I_p contradicting the fact that we are adding v_{ir}, or every petal of the sunflower contains a vertex in I_p, which is not possible since the size of any independent set is at most k and the number of petals is larger. Hence, we first modify the subsequence σ_s by adding v_{fr} instead of v_{ir}. Formally, we have $\sigma'_s = \langle (I_p \setminus \{v_{ir}\}) \cup \{v_{fr}\}, \ldots, (I_q \setminus \{v_{ir}\}) \cup \{v_{fr}\} \rangle$.

To be able to replace σ_s by σ'_s in σ and obtain a reconfiguration sequence from I_s to I_t, then all of the following conditions must hold:

(1) $|(I_q \setminus \{v_{ir}\}) \cup \{v_{fr}\}| = k$.
(2) $(I_i \setminus \{v_{ir}\}) \cup \{v_{fr}\}$ is an independent set of G for all $p \leq i \leq q$,
(3) $|(I_i \setminus \{v_{ir}\}) \cup \{v_{fr}\} \Delta (I_{i+1} \setminus \{v_{ir}\}) \cup \{v_{fr}\}| = 1$ for all $p \leq i < q$, and
(4) $k - 1 \leq |(I_i \setminus \{v_{ir}\}) \cup \{v_{fr}\}| \leq k$ for all $p \leq i \leq q$.

It is not hard to see that if there exists no i, $p < i \leq q$, such that σ'_s adds a vertex in $N[v_{fr}]$ at position i, then all four conditions hold. If there exists such a position, we will modify σ'_s into yet another subsequence σ''_s by finding a new vertex to represent v_{ir}. The length of σ''_s will be one greater than the length of σ'_s.

We let i, $p < i \leq q$, be the first position in σ'_s at which a vertex in $u \in N[v_{fr}]$ (possibly equal to v_{fr}) is added. Using the same arguments discussed to find v_{fr}, and since we constructed a sunflower with $2k$ petals, we can find another vertex v'_{fr} such that $N[v_{fr}] \cap I_{i-1} = \emptyset$. This new vertex will represent v_{ir} instead of v_{fr}. We construct σ''_s from σ'_s as follows: $\sigma''_s = \langle I_p \setminus \{v_{ir}\} \cup \{v_{fr}\}, \ldots, I_{i-1} \setminus \{v_{ir}\} \cup \{v_{fr}\}, I_{i-1} \setminus \{v_{ir}\} \cup \{v'_{fr}\}, I_i \setminus \{v_{ir}\} \cup \{v'_{fr}\}, \ldots, I_q \setminus \{v_{ir}\} \cup \{v'_{fr}\} \rangle$. If σ''_s now satisfies all four conditions then we are done. Otherwise, we repeat the same process (which can occur at most $q - p$ times) until we reach such a subsequence. □

Theorem 2. ISR *on d-degenerate graphs is fixed-parameter tractable parameterized by $k + d$.*

Proof. For an instance (G, I_s, I_t, k) of ISR, we know from Lemma 1 that as long as $V(G) \setminus \{I_s \cup I_t\}$ contains more than $(2d+1)!(2k-1)^{2d+1}$ vertices of degree at most $2d$ we can find an irrelevant vertex and reduce the size of the graph. After exhaustively reducing the graph to obtain G', we known that $G'[V(G') \setminus \{I_s \cup I_t\}]$, which is also d-degenerate, has at most $(2d+1)!(2k-1)^{2d+1}$ vertices of degree at most $2d$. Hence, applying Proposition 3, we know that $|V(G') \setminus \{I_s \cup I_t\}| \leq (2d+1)(2d+1)!(2k-1)^{2d+1}$ and $|V(G')| \leq (2d+1)(2d+1)!(2k-1)^{2d+1} + 2k$. □

3.2 Nowhere-dense graphs

Nesetril and Ossona de Mendez [37] showed an interesting relationship between nowhere-dense classes and a property of classes of structures introduced by Dawar [12, 13] called *quasi-wideness*. We will use quasi-wideness and show a rather interesting relationship between ISR on graphs of bounded degeneracy and nowhere-dense graphs. That is, our algorithm for nowhere-dense graphs will closely mimic the previous

algorithm in the following sense. Instead of using the sunflower lemma to find a large sunflower, we will use quasi-wideness to find a "large enough almost sunflower" with an initially "unknown" core and then use structural properties of the graph to find this core and complete the sunflower. We first state some of the results that we need. Given a graph G, a set $S \subseteq V(G)$ is called *r-scattered* if $N_G^r(u) \cap N_G^r(v) = \emptyset$ for all distinct $u, v \in S$.

Proposition 4. *Let G be a graph and let $S = \{s_1, s_2, ..., s_k\} \subseteq V(G)$ be a 2-scattered set of size k in G. Then the closed neighborhoods of the vertices in S form a sunflower with k petals and an empty core.*

Definition 5. *A class $\mathcal{C}$ of graphs is* uniformly quasi-wide *with margin $s_{\mathcal{C}} : \mathbb{N} \to \mathbb{N}$ and $N_{\mathcal{C}} : \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ if for all $r, k \in \mathbb{N}$, if $G \in \mathcal{C}$ and $W \subseteq V(G)$ with $|W| > N_{\mathcal{C}}(r, k)$, then there is a set $S \subseteq W$ with $|S| < s_{\mathcal{C}}(r)$, such that W contains an r-scattered set of size at least k in $G[V(G) \setminus S]$. $\mathcal{C}$ is* effectively uniformly quasi-wide *if $s_{\mathcal{C}}(r)$ and $N_{\mathcal{C}}(r, k)$ are computable.*

Examples of effectively uniformly quasi-wide classes include graphs of bounded degree with margin 1 and H-minor-free graphs with margin $|V(H)| - 1$.

Theorem 3 ([14]). *A class $\mathcal{C}$ of graphs is effectively nowhere-dense if and only if $\mathcal{C}$ is effectively uniformly quasi-wide.*

Theorem 4 ([14]). *Let $\mathcal{C}$ be an effectively nowhere-dense class of graphs and h be the computable function such that $K_{h(r)} \not\preceq_m^r G$ for all $G \in \mathcal{C}$. Let G be an n-vertex graph in $\mathcal{C}$, $r, k \in \mathbb{N}$, and $W \subseteq V(G)$ with $|W| \geq N(h(r), r, k)$, for some computable function N. Then in $\mathcal{O}(n^2)$ time, we can compute a set $B \subseteq V(G)$, $|B| \leq h(r) - 2$, and a set $A \subseteq W$ such that $|A| \geq k$ and A is an r-scattered set in $G[V(G) \setminus B]$.*

Lemma 2. *Let $\mathcal{C}$ be an effectively nowhere-dense class of graphs and h be the computable function such that $K_{h(r)} \not\preceq_m^r G$ for all $G \in \mathcal{C}$. Let (G, I_s, I_t, k) be an instance of ISR where $G \in \mathcal{C}$ and let R be the set of vertices in $V(G) \setminus \{I_s \cup I_t\}$. Moreover, let $\mathcal{P} = \{P_1, P_2, \ldots\}$ be a family of sets which partitions R such that for any two distinct vertices $u, v \in R$, $u, v \in P_i$ if and only if $N_G(u) \cap \{I_s \cup I_t\} = N_G(v) \cap \{I_s \cup I_t\}$. If there exists a set $P_i \in \mathcal{P}$ such that $|P_i| > N(h(2), 2, 2^{h(2)+1}k)$, for some computable function N, then there exists an irrelevant vertex $v \in V(G) \setminus \{I_s \cup I_t\}$ such that (G, I_s, I_t, k) is a yes-instance if and only if (G', I_s, I_t, k) is a yes-instance, where G' is obtained from G by deleting v and all edges incident on v.*

Proof. By construction, we known that the family $\mathcal{P}$ contains at most 4^k sets, as we partition R based on their neighborhoods in $I_s \cup I_t$. Note that some vertices in R have no neighbors in $I_s \cup I_t$ and will therefore belong to the same set in $\mathcal{P}$.

Assume that there exists a $P \in \mathcal{P}$ such that $|P| > N(h(2), 2, 2^{h(2)+1}k)$. Consider the graph $G[R]$. By Theorem 4, we can, in $\mathcal{O}(|R|^2)$ time, compute a set $B \subseteq R$, $|B| \leq h(2) - 2$, and a set $A \subseteq P$ such that $|A| \geq 2^{h(2)+1}k$ and A is a 2-scattered set in $G[R \setminus B]$. Now let $\mathcal{P}' = \{P_1', P_2', \ldots\}$ be a family of sets which partitions A such that for any two distinct vertices $u, v \in A$, $u, v \in P_i'$ if and only if $N_G(u) \cap B = N_G(v) \cap B$. Since $|A| \geq 2^{h(2)+1}k$ and $|\mathcal{P}'| \leq 2^{h(2)}$, we know that at least one set in $\mathcal{P}'$ will contain at least $2k$ vertices of A. Denote these $2k$ vertices by A'. All vertices in A' have the same neighborhood in B and the same neighborhood in $I_s \cup I_t$ (as all vertices in A' belonged to the same set $P \in \mathcal{P}$). Moreover, A' is a 2-scattered set in $G[R \setminus B]$. Hence, the sets $\{N_G[a_1'], N_G[a_2'], \ldots, N_G[a_{2k}']\}$, i.e. the closed neighborhoods of the vertices in A', form a sunflower with $2k$ petals (Proposition 4); the core of this sunflower is contained in $B \cup I_s \cup I_t$. Using the same arguments as we did in the proof of Lemma 1, we can show that there exists at least one irrelevant vertex $v \in V(G) \setminus \{B \cup I_s \cup I_t\}$. $\square$

Theorem 5. *ISR restricted to any effectively nowhere-dense class $\mathcal{C}$ of graphs is fixed-parameter tractable parameterized by k.*

Proof. If after partitioning $V(G) \setminus \{I_s \cup I_t\}$ into at most 4^k sets the size of every set $P \in \mathcal{P}$ is bounded by $N(h(2), 2, 2^{h(2)+1}k)$, then we can solve the problem by exhaustive enumeration, as $|V(G)| \leq 2k + 4^k N(h(2), 2, 2^{h(2)+1}k)$. Otherwise, we can apply Lemma 2 and reduce the size of the graph in polynomial time. $\square$

4 Dominating set reconfiguration

4.1 `W[1]`-hardness

The `W[1]`-hardness of the DSR problem can be shown using only minor modifications to the standard parameterized reduction from IS to DS. That is, instead of reducing from IS to DS, we can instead give a reduction from ISR to DSR. We include a proof for completeness.

Theorem 6. DSR *parameterized by k is* `W[1]`*-hard on general graphs.*

Proof. We let (G, I_s, I_t, k) be an instance of ISR, where $V(G) = \{v_1, \ldots, v_n\}$, $E(G) = \{e_1, \ldots, e_m\}$, $I_s = \{v_{i_1}, \ldots, v_{i_k}\}$, and $I_t = \{v_{j_1}, \ldots, v_{j_k}\}$. We first construct a graph G' as follows. G' consists of the disjoint union of k vertex-disjoint cliques $C_1, \ldots, C_k$, each of size n, k vertex-disjoint independent sets $F_1, \ldots, F_k$, each of size at most $k + 2$, and at most $n^2 k^2$ vertex-disjoint independent sets $R_1, R_2, \ldots$, each of size $k + 2$. Intuitively, each set F_i will force any dominating set of G' of size k (or $k + 1$) to pick a vertex from each C_i and the "R sets" will guarantee that the selected vertices form an independent set in G. Formally, we have:

(1) For every vertex $v \in V(G)$ there is a corresponding vertex in each C_i, $1 \leq i \leq k$ and we let $C_i = \{c_1^i, \ldots, c_n^i\}$.
(2) For every $1 \leq i \leq k$, we make the set C_i a clique in G'.
(3) For each set C_i, $1 \leq i \leq k$, we introduce a set F_i of $k + 2$ new independent vertices and add an edge between each vertex in C_i and all vertices in F_i.
(4) For a vertex $c_p^i \in C_i$ and a vertex $c_q^j \in C_j$, $i \neq j$, $1 \leq i, j \leq k$, and $1 \leq p, q \leq n$, if $p = q$ or $v_p v_q \in E(G)$ we introduce $k + 2$ new independent vertices and make them adjacent to all vertices in $C_i \cup C_j \setminus \{c_p^i, c_q^j\}$. In other words, each new vertex dominates all but two vertices in $C_i \cup C_j$, namely c_p^i and c_q^j.

We let (G', D_s, D_t, k) denote the corresponding DSR instance, where $D_s = \{c_{i_1}^1, \ldots, c_{i_k}^k\}$ and $D_t = \{c_{j_1}^1, \ldots, c_{j_k}^k\}$. Clearly, any dominating set D of G' of size k must pick exactly one vertex from each C_i, $1 \leq i \leq k$, and each such set corresponds to an independent set of size k in G. Moreover, any reconfiguration sequence between D_s and D_t starts by adding a vertex (since G' has no dominating set of size $k - 1$) and then removing another (since dominating sets larger than $k + 1$ are not allowed). By swapping the order of consecutive vertex additions and removals we obtain a one-to-one correspondence between reconfiguration sequences of independent sets of G (of size k and $k - 1$) and reconfiguration sequences (of the same length) between dominating sets of G' (of size k and $k + 1$). The instances are thus equivalent. □

4.2 Graphs excluding $K_{d,d}$ as a subgraph

The parameterized complexity of the DOMINATING SET problem (parameterized by k) on various classes of graphs has been studied extensively in the literature; the main goal has been to push the tractability frontier as far as possible. The problem was shown fixed-parameter tractable on planar graphs by Alber et al. [1], on bounded genus graphs by Ellis et al. [19], on H-minor-free graphs by Demaine et al. [15], on bounded expansion graphs by Nesetril and Ossona de Mendez [36], on nowhere-dense graphs by Dawar and Kreutzer [14], on degenerate graphs by Alon and Gutner [2], and finally on $K_{d,d}$-free graphs by Philip et al. [40] and Telle and Villanger [42]. Figure 1 illustrates the inclusion relationship among these classes of graphs, which all fall under the category of sparse graphs. Our fixed-parameter tractable algorithm relies on many of these earlier results. Interestingly, and since the class of $K_{d,d}$-free graphs includes all those other graph classes, our algorithm (Theorem 9) implies that the diameter of the reconfiguration graph $R_{\text{DS}}(G, k, k + 1)$ (or of its connected components), for G in any of the aforementioned classes, is bounded above by $f(k, c)$, where f is a computable function and c is constant which depends on the graph class at hand. We start with some definitions and known results.

Definition 6 ([18, 41, 40, 42]). *Given a graph G, the* domination core *of G is a set $C \subseteq V(G)$ such that any set $D \subseteq V(G)$ is a dominating set of G if and only if D dominates C. In other words, D is a dominating set of G if and only if $C \subseteq N_G[D]$.*

9

Theorem 7 ([41, 40, 42]). *If G is a graph which excludes $K_{d,d}$ as a subgraph and G has a dominating set of size at most k then the size of the domination core C of G is at most dk^d and C can be computed in $\mathcal{O}^*(dk^d)$ time.*

Definition 7. *A bipartite graph G with bipartition (A, B) is B-twinless if there are no vertices $u, v \in B$ such that $N(u) = N(v)$.*

Theorem 8 ([41]). *If G is a bipartite graph with bipartition (A, B) such that G is B-twinless and excludes $K_{d,d}$ as a subgraph then*

$$|B| \le 2(d-1)(\frac{|A|e}{d})^{2d}.$$

Since Theorem 7 implies a bound on the size of the domination core and allows us to compute it efficiently, our main concern is to deal with vertices outside of the core, i.e. vertices in $V(G) \setminus C$. The next lemma shows that we can in fact find strongly irrelevant vertices outside of the domination core of a graph.

Lemma 3. *For G an n-vertex graph, C the domination core of G, and D_s and D_t two dominating sets of G, if there exist $u, v \in V(G) \setminus \{C \cup D_s \cup D_t\}$ such that $N_G(u) \cap C = N_G(v) \cap C$ then u (or v) is strongly irrelevant.*

Proof. Given a reconfiguration sequence $\sigma = \langle D_0 = D_s, D_1, \dots, D_\ell = D_t \rangle$ from D_s to D_t which touches u, we will show how to obtain a reconfiguration sequence σ' such that $|\sigma'| \le |\sigma|$ and σ' touches v but not u.

We construct σ' in two stages. In the first stage, we construct the sequence $\alpha = \langle D_0', D_1', \dots, D_\ell' \rangle$ of dominating sets, where for all $0 \le i \le \ell$

$$D_i' = \begin{cases} D_i \cup \{v\} \setminus \{u\} \text{ if } u \in D_i \\ D_i \text{ if } u \notin D_i. \end{cases}$$

Note that α is not necessarily a reconfiguration sequence from D_s to D_t. In the second stage, we repeatedly delete from α any set D_i' such that $D_i' = D_{i+1}'$, $0 \le i < \ell$. We let $\sigma' = \langle D_0', D_1', \dots, D_{\ell'}' \rangle$ denote the resulting sequence, in which there are no two consecutive sets that are equal, and we claim that σ' is in fact a reconfiguration sequence from D_s to D_t.

To prove the claim, we need to show that the following conditions hold:

(1) $D_0' = D_s$ and $D_{\ell'}' = D_t$,
(2) D_i' is a dominating set of G for all $0 \le i \le \ell'$,
(3) $|D_i' \Delta D_{i+1}'| = 1$ for all $0 \le i < \ell'$, and
(4) $k \le |D_i'| \le k + 1$ for all $0 \le i \le \ell'$.

Since $u, v \notin D_s \cup D_t$, condition (1) clearly holds. Moreover, since replacing u by v in any set does not increase the size of the corresponding set, $k \le |D_i'| \le k + 1$ (condition (4) holds) and $|D_i' \Delta D_{i+1}'| \le 1$. As there are no two consecutive sets in σ' that are equal, $|D_i' \Delta D_{i+1}'| > 0$ and therefore $|D_i' \Delta D_{i+1}'| = 1$ (condition (3) holds). The fact that D_i' is a dominating set of G follows from the definition of a domination core. Since D_i is a dominating set of G, $C \subseteq N_G[D_i]$. Moreover, since $N_G(u) \cap C = N_G(v) \cap C$ and $u, v \notin C$, we know that $C \subseteq N_G[D_i']$. By the definition of the domination core, it follows that D_i' (which still dominates C) is also a dominating set of G. Therefore, all four conditions hold, as needed. $\square$

Theorem 9. DSR *parameterized by $k + d$ is fixed-parameter tractable on graphs that exclude $K_{d,d}$ as a subgraph.*

Proof. Given a graph G, integer k, and two dominating sets D_s and D_t of G of size at most k, we first compute the domination core C of G, which by Theorem 7 can be accomplished in $\mathcal{O}^*(dk^d)$ time. Next, and due to Lemma 3, we can delete all strongly irrelevant vertices from $V(G) \setminus \{C \cup D_s \cup D_t\}$. We denote this new graph by G'.

10

Now consider the bipartite graph G'' with bipartition $(A = C \backslash \{D_s \cup D_t\}, B = V(G') \backslash \{C \cup D_s \cup D_t\})$. This graph is B-twinless, since for every pair of vertices $u, v \in V(G) \backslash \{C \cup D_s \cup D_t\}$ such that $N_G(u) \cap C = N_G(v) \cap C$ either u or v is strongly irrelevant and is therefore not in $V(G')$ nor $V(G'')$. Moreover, since every subgraph of a $K_{d,d}$-free graph is also $K_{d,d}$-free, G'' is $K_{d,d}$-free. Hence, by Theorems 7 and 8, we have

$$|B| \le 2(d-1)(\frac{|A|e}{d})^{2d}$$
$$\le 2d(3|A|)^{2d} \le 2d(3dk^d)^{2d}.$$

Putting it all together, we know that after deleting all strongly irrelevant vertices, the number of vertices in the resulting graph G' is at most

$$|V(G')| = |V(C)| + |D_s \cup D_t| + |V(G') \setminus \{C \cup D_s \cup D_t\}|$$
$$\le dk^d + 2k + 2d(3dk^d)^{2d}$$

Hence, we can solve DSR by exhaustively enumerating all $2^{|V(G')|}$ subsets of $V(G')$ and building the reconfiguration graph $R_{\mathrm{DS}}(G', k, k+1)$. □

References

1. J. Alber, H. L. Bodlaender, H. Fernau, and R. Niedermeier. Fixed parameter algorithms for planar dominating set and related problems. In *Proceedings of the 17^{th} Scandinavian Workshop on Algorithm Theory*, pages 97–110. Springer Berlin Heidelberg, 2000.
2. N. Alon and S. Gutner. Linear time algorithms for finding a dominating set of fixed size in degenerated graphs. *Algorithmica*, 54(4):544–556, 2009.
3. M. Bonamy and N. Bousquet. Recoloring bounded treewidth graphs. *Electronic Notes in Discrete Mathematics*, 44:257–262, 2013.
4. P. Bonsma. The complexity of rerouting shortest paths. In *Proceedings of the 37^{th} International Symposium on Mathematical Foundations of Computer Science*, pages 222–233, 2012.
5. P. Bonsma. Rerouting shortest paths in planar graphs. In *Proceedings of the 32^{nd} Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, pages 337–349, 2012.
6. P. Bonsma and L. Cereceda. Finding paths between graph colourings: PSPACE-completeness and superpolynomial distances. *Theoretical Computer Science*, 410(50):5215–5226, 2009.
7. P. Bonsma, A. E. Mouawad, N. Nishimura, and V. Raman. The complexity of bounded length graph recoloring and CSP reconfiguration. In *Proceedings of the 9^{th} International Symposium on Parameterized and Exact Computation*, pages 110–121, 2014.
8. A. Brandstädt, V. B. Le, and J. P. Spinrad. *Graph Classes: A Survey*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1999.
9. L. Cereceda. *Mixing graph colourings*. PhD thesis, London School of Economics, 2007.
10. L. Cereceda, J. van den Heuvel, and M. Johnson. Mixing 3-colourings in bipartite graphs. *European Journal of Combinatorics*, 30(7):1593–1606, 2009.
11. L. Cereceda, J. van den Heuvel, and M. Johnson. Finding paths between 3-colorings. *Journal of Graph Theory*, 67(1):69–82, 2011.
12. A. Dawar. Finite model theory on tame classes of structures. In *Proceedings of the 32^{nd} International Symposium on Mathematical Foundations of Computer Science*, volume 4708 of *Lecture Notes in Computer Science*, pages 2–12. Springer Berlin Heidelberg, 2007.
13. A. Dawar. Homomorphism preservation on quasi-wide classes. *Journal of Computer and System Sciences*, 76(5):324–332, 2010. Workshop on Logic, Language, Information and Computation.
14. A. Dawar and S. Kreutzer. Domination problems in nowhere-dense classes of graphs. *CoRR*, 2009. arXiv:0907.42837.
15. E. D. Demaine, F. V. Fomin, M. Hajiaghayi, and D. M. Thilikos. Subexponential parameterized algorithms on bounded-genus graphs and H-minor-free graphs. *Journal of the ACM*, 52(6):866–893, 2005.
16. R. Diestel. *Graph theory*. Springer-Verlag, Electronic Edition, 2005.
17. R. G. Downey and M. R. Fellows. *Parameterized complexity*. Springer-Verlag, New York, 1997.

18. P. G. Drange, M. S. Dregi, F. V. Fomin, S. Kreutzer, D. Lokshtanov, M. Pilipczuk, M. Pilipczuk, F. Reidl, S. Saurabh, F. S. Villaamil, and S. Sikdar. Kernelization and sparseness: the case of dominating set. *CoRR*, 2014. arXiv:1411.4575.

19. J. Ellis, H. Fan, and M. Fellows. The dominating set problem is fixed parameter tractable for graphs of bounded genus. In *Proceedings of the 19^{th} Scandinavian Workshop on Algorithm Theory*, volume 2368 of *Lecture Notes in Computer Science*, pages 180–189. Springer Berlin Heidelberg, 2002.

20. P. Erdos and R. Rado. Intersection theorems for systems of sets. *Journal of the London Mathematical Society*, 35:85–90, 1960.

21. J. Flum and M. Grohe. *Parameterized complexity theory*. Springer-Verlag, Berlin, 2006.

22. P. Gopalan, P. G. Kolaitis, E. N. Maneva, and C. H. Papadimitriou. The connectivity of Boolean satisfiability: computational and structural dichotomies. *SIAM Journal on Computing*, 38(6):2330–2355, 2009.

23. M. Grohe, S. Kreutzer, and S. Siebertz. Characterisations of nowhere dense graphs. In *Proceedings of the 33^{rd} Annual Conference on Foundations of Software Technology and Theoretical Computer Science*, pages 21–40, 2013.

24. M. Grohe, S. Kreutzer, and S. Siebertz. Deciding first-order properties of nowhere dense graphs. In *Proceedings of the 46^{th} Annual ACM Symposium on Theory of Computing*, pages 89–98, 2014.

25. T. Ito, E. D. Demaine, N. J. A. Harvey, C. H. Papadimitriou, M. Sideri, R. Uehara, and Y. Uno. On the complexity of reconfiguration problems. *Theoretical Computer Science*, 412(12-14):1054–1065, 2011.

26. T. Ito, M. Kamiński, and E. D. Demaine. Reconfiguration of list edge-colorings in a graph. *Discrete Applied Mathematics*, 160(15):2199–2207, 2012.

27. T. Ito, M. Kamiski, and H. Ono. Fixed-parameter tractability of token jumping on planar graphs. In *Algorithms and Computation*, Lecture Notes in Computer Science, pages 208–219. Springer International Publishing, 2014.

28. T. Ito, M. Kamiski, H. Ono, A. Suzuki, R. Uehara, and K. Yamanaka. On the parameterized complexity for token jumping on graphs. In *Proceedings of the 11^{th} Annual Conference on Theory and Applications of Models of Computation*, volume 8402 of *Lecture Notes in Computer Science*, pages 341–351. Springer International Publishing, 2014.

29. M. Johnson, D. Kratsch, S. Kratsch, V. Patel, and D. Paulusma. Finding shortest paths between graph colourings. *CoRR*, 2014. arXiv:1403.6347.

30. M. Kamiński, P. Medvedev, and M. Milanič. Shortest paths between shortest paths. *Theoretical Computer Science*, 412(39):5205–5210, 2011.

31. M. Kamiński, P. Medvedev, and M. Milanič. Complexity of independent set reconfigurability problems. *Theoretical Computer Science*, 439:9–15, 2012.

32. D. R. Lick and A. T. White. k-degenerate graphs. *Canadian Journal of Mathematics*, 22:1082–1096, 1970.

33. A. E. Mouawad, N. Nishimura, and V. Raman. Vertex cover reconfiguration and beyond. *CoRR*, 2014. arXiv:1402.4926.

34. A. E. Mouawad, N. Nishimura, V. Raman, N. Simjour, and A. Suzuki. On the parameterized complexity of reconfiguration problems. In *Proceedings of the 8^{th} International Symposium on Parameterized and Exact Computation*, 2013.

35. A. E. Mouawad, N. Nishimura, V. Raman, and M. Wrochna. Reconfiguration over tree decompositions. *CoRR*, 2014. arXiv:1405.2447.

36. J. Nesetril and P. O. de Mendez. Structural properties of sparse graphs. In *Building Bridges*, volume 19 of *Bolyai Society Mathematical Studies*, pages 369–426. Springer Berlin Heidelberg, 2008.

37. J. Nesetril and P. O. de Mendez. First order properties on nowhere dense structures. *Journal of Symbolic Logic*, 75(3):868–887, 2010.

38. J. Nesetril and P. O. de Mendez. From sparse graphs to nowhere dense structures: Decompositions, independence, dualities and limits, 2010. European Congress of Mathematics.

39. R. Niedermeier. *Invitation to fixed-parameter algorithms*. Oxford University Press, Oxford, 2006.

40. G. Philip, V. Raman, and S. Sikdar. Solving dominating set in larger classes of graphs: FPT algorithms and polynomial kernels. In *Proceedings of the 17^{th} Annual European Symposium on Algorithms*, volume 5757 of *Lecture Notes in Computer Science*, pages 694–705. Springer Berlin Heidelberg, 2009.

41. S. Saurabh. Private communications, 2014.

42. J. A. Telle and Y. Villanger. FPT algorithms for domination in biclique-free graphs. In *Proceedings of the 20^{th} Annual European Conference on Algorithms*, pages 802–812, 2012.

43. J. van den Heuvel. The complexity of change. *Surveys in Combinatorics 2013*, 409:127–160, 2013.

44. M. Wrochna. Reconfiguration in bounded bandwidth and treedepth. *CoRR*, 2014. arXiv:1405.0847.